Eco514—Game Theory
Lecture 8: Applications (1)—Simultaneous Auctions

Marciano Siniscalchi

October 12, 1999

Introduction

This lecture, as well as the next, exemplify applications of the framework and techniques developed so far to problems of economic interest. Neither lecture attempts to cover the example applications in any generality, of course; you may however find these topics of sufficient interest to warrant further study.

Auction theory is generally indicated as one of the "success stories" of game theory. There is no doubt that the game-theoretic analysis of auctions has informed design decisions of great practical relevance. If you are interested in recent applications of auction theory, you should definitely check out the following URL: `http://www.market-design.com` (in particular, the Library section contains downloadable working papers in PDF format).

This lecture focuses on the simplest possible setting, in which a single good is offered for sale, and bidders are required to submit their offers in a sealed envelope (so we are ruling out ascending or descending auctions: but see below). Also, I shall focus on the analysis of "popular" auction forms: Myerson's 1981 paper deals with optimal auctions from a mechanism design standpoint. Finally, I will employ equilibrium analysis throughout.

Framework and Auction Forms

The basic bidding game form may be represented as follows: each bidder $i \in N$ submits a bid $a_i \in A_i = \mathbf{R}_+$ (so we rule out negative bids); given any bid profile $a = (a_i)_{i \in N}$, the auctioneer determines the set of *high bidders* $H(a)$ according to

$$H(a) = \{i \in N : \forall j \neq i,\ a_i \geq a_j\}.$$

Finally, the object is randomly assigned to one of the high bidders (each of whom is equally likely to receive it), referred to as the *winning bidder*. Thus, the winning bidder is a high bidder, but only one high bidder will be the winning bidder.[1]

VISIT…

The bulk of the literature assumes quasi-linear preferences (although risk-aversion does play a role, and is analyzed in several papers). However, this is clearly not sufficient to determine payoffs. We still need to specify: (i) how much the object is worth to each bidder (we shall assume quasi-linear utility); and (ii) the bidders' transfers to the seller.[2] Moreover, unless we assume that players' valuations are known, we need to specify (iii) the form of payoff uncertainty we are interested in.

Specifying (i), (ii) and (iii) leads to a taxonomy of bidding games:

- With respect to (i), we distinguish between *private* and *interdependent* values: in the former setting, each player knows her valuation for the object (but may be uncertain about her opponents' valuation); in the latter, one or more players may be uncertain about her valuation, which will typically be related to other players' valuations and/or realizations of underlying state variables. In particular, if the object is worth the same to all bidders, but no bidder is certain of the value, we are in a *pure common values* setting.

- Payment rules (ii) pin down the auction mechanism, the most popular being *first-price*, *second-price* and (to a lesser degree) all-pay auctions.

- Finally, with respect to (iii), whatever information bidders receive or hold (be it their valuation or a *signal* thereof) may be either *independent* or *correlated*.

A few examples will clarify (i) and (iii): the differences may appear to be subtle, but they are substantial.

Suppose that each bidder $i \in N$ observes a *signal* $s_i \in [0,1]$, that F denotes the joint cdf of the signals, and that bidder i's valuation is a function $v_i : [0,1]^N \to \mathbf{R}_+$. This is a very common specification of payoff uncertainty.

In this simplified (but very popular) framework, (i) and (iii) entail restrictions on the functions v_i and F respectively. If $v_i(s_i, s_{-i}) = v_i(s_i, s'_{-i})$ for all $s_i \in [0,1]$ and $s_{-i} \in [0,1]^{N-1}$, then we are in a setting with private values: bidder i's valuation is uniquely determined by her own signal s_i, which she observes. In the simplest case, $v_i(s_i, s_{-i}) = s_i$ (which, note well, also entails a symmetry assumption).

[1]Sellers may also set a *reserve price*: all bids below the reserve price are simply rejected. In this case, the object may remain unsold if all bids are below the reserve price. Although Myerson shows that, in certain situations, setting a positive reserve price maximizes expected revenues, we shall simply ignore this possibility.

[2]In the most common auction formats, only the winning bidder is required to pay. However, in *all-pay* auctions, as the name suggests, all bidders pay their bid. Indeed, one can conjure up just about any kind of payment rule: as long as the winning bidder is a high bidder, you can still call the resulting mechanism an "auction".

Otherwise, valuations are interdependent; for instance, a pedagogically useful and convenient specification is $v_i(s) = \sum_{i \in N} s_i$. This is known as the "wallet game[3]" (note again the implicit symmetry assumption).

The latter is a specification with pure common values. The expression "interdependent values" allows for intermediate cases such as $v_i(s) = a_i s_i + \sum_{j \neq i} s_j$, for some $a_i > 1$.

In either case, the signals s_i, $i \in N$, may be independent (so that $F(s) = \prod_{i \in N} F_i(s_i)$ for some collection of one-dimensional cdf's F_i, $i \in N$) or correlated; for instance, suppose that there exist $N + 1$ i.i.d. random variables $x_0, x_1, \ldots, x_N$, uniform on $[0, \frac{1}{2}]$, such that $s_i = x_0 + x_i$ for every $i \in N$.

The key intuition is that, since bids will be functions of the signals, when signals are correlated, bidder i can make inferences about her opponents' equilibrium bids based on her own signal. This is of course true regardless of whether values are private or interdependent.

Finally, payment rules are relatively simple to specify. In a *first-price auction*, the winning bidder pays her bid. Thus, the *payment rule* $m : A \to \mathbf{R}_+^N$ is defined by

$$m_i(a) = \begin{cases} a_i & i \in H(a) \\ 0 & i \notin H(a) \end{cases}$$

In a *second-price auction*, she pays the *highest non-winning bid*: that is,

$$m_i(a) = \begin{cases} \max_{j \neq i} a_j & i \in H(a) \\ 0 & i \notin H(a) \end{cases}$$

Auctions are modelled as (infinite) Bayesian games. In the simplified setting we are looking at, signals contain all payoff-relevant information (although other formulations could be considered). It is typical to let

$$\Omega = S \quad \text{and} \quad t_i(s) = s_i \times [0,1]^{N-1}$$

reflecting the assumption that players know their signal, and nothing else. Indeed, in this setting, the distinction between the signal s_i and the corresponding type $s_i \times [0,1]^{N-1}$ is blurred, and I will follow conventional usage by adopting whatever term is most convenient.

The cdf F is then used to define a common prior on Ω. As I have already noted, this has several implications for equilibrium analysis; one could entertain different hypotheses, but this is not the point of today's lecture.

Also note that the resulting game $G = (N, \Omega, (A_i, u_i, t_i)_{i \in N})$ is infinite both because action spaces are $(A_i = \mathbf{R}_+)$ and because the state space is. It should be noted that bidding

[3]The idea is as follows: players place their wallets on the table, then bid for the money contained in all wallets. Each participant knows how much money she is carrying in her own wallet, but ignores the content of the other wallets.

games exhibit serious discontinuities, so that the existence of Nash equilibria, and even best replies (see below) is not guaranteed by the standard machinery.

Quasi-linearity completes the specification of payoffs: fixing a payment rule m, for any profile $a \in A$,

$$u_i(a,s) = \begin{cases} \frac{1}{|H(a)|}[v_i(s) - m_i(a)] & i \in H(a) \\ 0 & i \notin H(a) \end{cases}$$

Analysis

We now consider a few even more special illustrative cases. Before we begin, consider the following definition:

Definition 1 *Fix a bidding game G defined as above, and a player $i \in N$. A* bid function *for bidder i is a function $\mathbf{a}_i : S_i \to A_i$.*

That is, a bid function is a degenerate signal-contingent randomized action.

Second-Price Auctions with Private Values

Let us assume for notational simplicity that $v_i(s) = s_i$. These bidding games are traditionally analyzed using the notion of *weak dominance*. An action a_i weakly dominates an alternative action a_i' for type s_i of bidder i iff, for all opponent bid functions $(\mathbf{a}_j)_{j \neq i}$ and for all realization of the opponents' signals,

$$u_i(a_i, (\mathbf{a}_j(s_j))_{j \neq i}, s_i, s_{-i}) \geq u_i(a_i', (\mathbf{a}_j(s_j))_{j \neq i}, s_i, s_{-i})$$

and moreover there exists a profile of opponents' bid functions and a profile of signal realizations for which the above inequality is strict.

With private values, however, this reduces to the requirement that, for all profiles of opponents' *actions* a_{-i} (not bid functions!),

$$u_i(a_i, a_{-i}, s_i, s_{-i}) \geq u_i(a_i', a_{-i}, s_i, s_{-i}) \quad \forall s_{-i} \in [0,1]^{N-1}$$

and that the inequality be strict for at least one profile of opponents' actions; this is a consequence of the assumption that s_{-i} does not affect $v_i(s)$.

But this leads to the conclusion that, for any signal realization $s \in [0,1]^N$, bidder $i \in N$ and pair of actions a_i, a_i', a_i weakly dominates a_i' for type s_i of bidder i iff a_i weakly dominates a_i' in the game *without payoff uncertainty* defined by $G(s) = \{N, (A_i, u_i(\cdot, s))_{i \in N}\}$. This is the game you analyzed in Problem Set 1, so you have already proved the following result:

Proposition 0.1 *In any second-price auctions with private values, bidding one's valuation is the unique weakly dominant action. Therefore, the profile of "truthful" bid functions defined by* $\mathbf{a}_i(s_i) = s_i$ *for every* $i \in N$*, is a Bayesian Nash equilibrium of the associated bidding game.*

Thus, truthful bidding is a Bayesian Nash equilibrium of the second-price auction—regardless of the assumptions one makes about the underlying signal structure. This is viewed as a particularly comforting result: truthful bidding is certainly "focal" in this setting; Proposition 0.1 shows that it follows from the assumption that players do not choose weakly dominated actions. You already know that the (complete-information version of the) game has other Bayesian Nash equilibria.

If participants bid truthfully, then no bidder has a higher valuation for the object than the winner. Thus, under truthful bidding, the second-price auction places the object in the "right" hands; that is, it is an *efficient* mechanism.[4]

Incidentally, you may already know that second-price auctions are special cases of Vickrey-Groves-Clark direct mechanisms. With private values, truthful reporting is weakly dominant in such mechanisms, and always leads to efficient outcomes. Thus, for example, the k-good extension of the second-price auction (known as the Vickrey auction) also achieves efficiency.

First-Price Auctions with Independent Private Values

The analysis of first-price auctions requires more work, and more restrictive assumptions. In particular, we need to assume a lot of symmetry: in particular, assume that $v_i(s_i, s_{-i}) = v(s_i)$ for all $s = (s_i, s_{-i}) \in [0,1]^N$ (so the same function $v : [0,1] \to \mathbf{R}_+$ defines valuations for all players) and $F_i = \bar{F}$ for all $i \in I$ (so signals are i.i.d.). Also assume that $v(1) = \bar{v} < \infty$ and $v(0) = 0$. Finally, we assume that v and $\bar{F}$ are continuously differentiable, and that the density $\bar{f} = \bar{F}'$ is bounded away from zero.

Now observe that, since the valuation function is bounded, we can renormalize it without loss of generality so that $v(1) = 1$. Having done this, since both v and $\bar{F}$ are continuous and v is increasing, we can further simplify both the notation and the algebra, again without loss of generality, by considering an equivalent setting in which $v(s) = s$ and $\bar{F}$ is replaced with a new cdf G such that the density of any valuation $x \in [0,1]$ under the original model, $\bar{f}(v^{-1}(x))$, equals the density of x under G, i.e. $g(x)$.

We look for a symmetric equilibrium in increasing, differentiable bid functions. The strategy we shall follow is to assume that such an equilibrium profile $(\mathbf{a}_i)_{i \in N} = (\mathbf{a}, \ldots, \mathbf{a})$ exists, and derive some conditions which the profile must satisfy in order to be an equilibrium in differentiable bid functions. Then, we show that any profile of differentiable bid functions satisfying the above conditions is indeed an equilibrium.

[4]Note well: in this literature, an auction mechanism is deemed efficient if there is *at least one* equilibrium of the associated game which induces an efficient allocation. Other equilibria may well be inefficient.

Note that this approach does *not* establish existence directly: first, no profile of bid functions satisfying the derived conditions may exist (so the second part of the argument may be empty); second, even if they do exist, they may fail to be increasing and/or differentiable (so the first part of the argument may be unwarranted). To establish existence, one must invoke results from the theory of differential equations.

With this in mind, note that Bidder i's payoff function when her opponents follow the equilibrium bid functions can be written as follows:

$$U_i(a_i, s_i) = \sum_{J \subset N \setminus \{i\}} \frac{1}{|J|+1} \Pr\left\{s_{-i} : H(a_i, (\mathbf{a}(s_j))_{j \neq i}) = J \cup \{i\}\right\} [s_i - a_i].$$

That is: if she wins the object by bidding a_i, Bidder i's utility is $s_i - a_i$; for any realization of the signals, she may be the only high bidder, or there may be another high bidder, or there may be two, etc., in which case she wins with correspondingly smaller probability.

The assumptions we have made simplify the analysis considerably. First, since the model is perfectly symmetric, we can drop all indices from actions and signals.

Second, since we are assuming that the common equilibrium bid function is increasing, and that G has a density, ties occur with probability zero (i.e. there is G-almost always only one winner). Thus, only one term (corresponding to $J = N \setminus \{i\}$) is nonzero in the above summation. Third, since the equilibrium bid function is increasing and signals are independent, it can be inverted, and the probability that the bid a is the winning bid is given by $G(\mathbf{a}^{-1}(a))^{N-1}$; this is the probability that all opponents observe signals which induce them (in the equilibrium under consideration) to bid at most a. Thus, we get

$$U(a, s) = [G(\mathbf{a}^{-1}(a))]^{N-1}[s - a]$$

Now, if a is a best reply to the equilibrium belief, then it must satisfy the first-order condition

$$(N-1)G(\mathbf{a}^{-1}(a))^{N-2} g(\mathbf{a}^{-1}(a)) \frac{1}{\mathbf{a}'(\mathbf{a}^{-1}(a))}[s - a] - [G(\mathbf{a}^{-1}(a))]^{N-1} = 0.$$

If the profile under consideration is an equilibrium then the above restriction must hold for $a = \mathbf{a}(s)$; this allows one to simplify the expression somewhat. Moreover, since G is strictly positive by assumption, we can divide throughout by $G(s)^{N-2}$ to get

$$(N-1)g(s)[s - \mathbf{a}(s)] = G(s)\mathbf{a}'(s); \tag{1}$$

the interpretation in terms of cost-benefit tradeoff should be clear. Another necessary condition for $\mathbf{a}$ to be a best-reply is $\mathbf{a}(0) = 0$. Thus, we have a first-order linear differential equation and a boundary condition; given our regularity conditions, a solution exists and is determined by the above necessary conditions.

Incidentally, this is where symmetry plays a role. Without it, Equation (1) would be replaced by a system of differential equations, and a more sophisticated machinery (or some clever trick) is required to establish the existence of a solution. We shall compute it explicitly under the additional assumption that $G(s) = s$, i.e. signals are uniformly distributed. Then the above differential equation reduces to

$$(N-1)[s - \mathbf{a}(s)] = s\mathbf{a}'(s)$$

whence we can verify immediately that $\mathbf{a}(s) = \frac{N-1}{N}s$ is the only solution satisfying $\mathbf{a}(0) = 0$. Given this parametrization, it is easy to verify that the profile $(\frac{N-1}{N}s, \ldots, \frac{N-1}{N}s)$ is indeed Bayesian Nash equilibrium of the first-price auction. Indeed, it is the *only* symmetric equilibrium in differentiable and increasing strategies. Note that this verification is necessary, because the differential equation derived above was only a necessary condition for the function $\mathbf{a}(\cdot)$ to be a best reply.

More generally, it can be shown that Equation (1), along with the boundary condition $\mathbf{a}(0) = 0$, indeed defines a Bayesian Nash equilibrium of the game.

In case you are wondering whether invoking, say, concavity of the objective function to conclude that Equation (1) is also a sufficient condition for $\mathbf{a}$ to define an equilibrium, think about the following: the concavity of the objective function $U(a, s)$ depends on exogenously given entities such as v and G, but also on the function $\mathbf{a}$ itself.

Two observations are in order. First, note that bidders *shade* their bids—the latter is always below their valuation, although the difference shrinks to zero in the limit. However, under the same assumptions about signal distributions and valuations, *first- and second-price auctions yield the same expected revenues* to the seller.

To see this in our simple example, note that at any state s, the price paid to the seller equals the second-highest signal; hence, expected revenues are equal to the expectation of the second order statistic drawn from an i.i.d. uniform sample of size N. This is

$$\int s \frac{N}{(N-2)!1!} G(s)^{N-2}[1 - G(s)]g(s)ds = N(N-1)\int s^{N-1}(1-s)ds = \frac{N-1}{N+1}.$$

On the other hand, expected revenues in a first-price auction equal $\frac{N-1}{N}$ times the expectation of the first order statistic,

$$\int sNG(s)^{N-1}g(s)ds = N\int s^N ds = \frac{N}{N+1}.$$

Again, this is actually an instance of a more general fact (see Myerson's paper for details): assuming we are willing to compare equilibria in different games, first- and second-price auctions are *revenue-equivalent* with private values.

Notes on interdependent values

I will limit my remarks to the simplest possible example of bidding in an interdependent values-regime: the wallet game. In particular, assume there are only two bidders.

Let us consider the second-price auction first. A crucial difference between private- and interdependent-values settings is that *players no longer have weakly dominant strategies.*

With private values, a player who bids less than her valuation risks losing the object, and gains nothing conditional upon winning because her payment is not determined by her bid, but by her opponent's.

With interdependent values, her valuation *conditional upon winning the object* depends on her opponent's bid function, because the latter determines the set of types against which she wins; thus, indirectly, her own bid determines her valuation conditional upon winning, and the entire logic of the Vickrey-Groves-Clark mechanism breaks down.

Indeed, "truthful bidding" is not even defined, if not perhaps with reference to a fixed equilibrium belief! Thus, we can at best hope to obtain efficiency in equilibrium, not as a consequence of the sole assumption that players do not choose weakly dominated actions. Indeed, we shall see that efficiency may fail to obtain in equilibrium.

Let us solve the wallet game under the second-price payment rule: recall that $v_i(s) = s_1 + s_2$. I claim that the profile $(2s_1, 2s_2)$ is a Bayesian Nash equilibrium of the associated game. We need only check that $2s$ is a best reply for type s against the putative equilibrium belief. Her expected payoff when she bids a is given by

$$U(a,s) = \int_0^{\frac{1}{2}a} (s + s_2 - 2s_2) ds_2 = \frac{1}{2}as - \frac{1}{2}a^2$$

which has a maximum at $a = 2s$, as one can easily check; the second-order conditions are also easily verified. Also note that, since values are purely common, any allocation mechanism is trivially efficient. However, there are examples of three-bidder games with interdependent values in which efficiency fails.

Note that revenues equal twice the expected value of the minimum of two i.i.d. uniform random variables, or (from our previous analysis) $\frac{2}{3}$.

Let us now consider a first-price auction in this setting. I claim than now (s_1, s_2) is an equilibrium profile. To see this, note that

$$U(a,s) = \int_0^{a} (s + s_2 - a) ds_2 = a(s-a) + \tfrac{1}{2}a^2$$

and again the claim is proved. From our previous analysis, expected revenues equal the expected value of the maximum of two i.i.d. uniform variates, or (again) $\frac{2}{3}$. This, however, is *not* generally the case for interdependent-values settings.

Micro-Notes on Correlation

The foregoing discussion leaves open the issue of correlation vs. independence of signals. I will merely point out the differences between these two cases in a limiting case: you should consult Myerson's 1981 paper, or Cremer and McLean's paper, for more on the topic.

Let us consider a two-bidder first-price auction with private values and *perfectly correlated signals*: that is, $\Pr\{s_1 = s_2\} = 1$. Thus, as soon as a bidder learns her signal, she also learns her opponent's signal; for all practical purposes, at the *interim* stage the players are actually engaged in a game with complete information.

It should be clear that the only equilibrium in the game where both signals equal some $s \in [0, 1]$ is for both players to bid exactly s, i.e. their valuation (if the winning bidder bids more than her opponent, she gains by reducing her bid; if both are bidding $a < s$, then either one can gain by raising her bid slightly above a). Thus, there is no shading, and the seller is able to extract the whole surplus from the players.

This turns out to be a general result. In any case, this simple example emphasizes that the stochastic structure of the signals matters.